

SmartSense Installation 1

## Installing SmartSense on HDF

**Date of Publish:** 2018-11-15



<http://docs.hortonworks.com>

# Contents

<b>SmartSense installation.....</b>	<b>3</b>
SmartSense system requirements.....	3
Operating system, JDK, and browser requirements.....	3
Software requirements.....	3
Ambari requirements.....	3
Installing SmartSense.....	3
Download and Install SmartSense.....	3
HST server placement.....	4
Add the SmartSense service.....	4
Installing SmartSense gateway.....	5
SmartSense gateway placement.....	6
Install and start SmartSense gateway.....	6
Integrating gateway with Ambari-managed SmartSense.....	7
<b>SmartSense uninstallation.....</b>	<b>7</b>
Uninstall SmartSense Gateway.....	7
<b>SmartSense ports and traffic flow.....</b>	<b>8</b>
User interface or Ambari View to HST server.....	8
HST agents to HST server.....	8
HST server to SmartSense gateway.....	9
SmartSense gateway to Hortonworks.....	9
<b>SmartSense installation troubleshooting.....</b>	<b>10</b>
SmartSense SSL troubleshooting.....	10
Reporting issues.....	11
<b>Enabling flex support subscription.....</b>	<b>11</b>

## SmartSense installation

Installing Hortonworks SmartSense Tool (HST) on a Hortonworks Data Flow (HDF) cluster involves meeting minimum system requirements, installing SmartSense, and choosing how bundles will be uploaded to Hortonworks.

### SmartSense system requirements

To run SmartSense, your system must meet the following requirements.

Your system must meet requirements in the following areas:

- Operating system requirements
- JDK requirements
- Browser requirements
- Software requirements
- Ambari requirements

### Operating system, JDK, and browser requirements

To learn about the operating system, JDK, and browser requirements for SmartSense, refer to the Support Matrix.

#### Related Information

[Support Matrix](#)

### Software requirements

You must install the following packages on each of the hosts in your cluster. These packages are used to gain a more complete diagnostic profile of the cluster.

- wget
- sysstat
- dstat
- lsof
- net-tools
- Python2, version 2.6 or later

### Ambari requirements

SmartSense is integrated with and deployed through Apache Ambari. SmartSense can be installed on an HDF cluster with Ambari 2.7.0 and later and it can be used with HDF 3.2.x or later.

## Installing SmartSense

Installing SmartSense involves installing SmartSense manually, reviewing HST server placement, and adding SmartSense service in Ambari.

### Download and Install SmartSense

Perform these steps to install SmartSense on your HDF cluster.

Steps

1. Obtain root permissions by running `sudo su`
2. Install SmartSense on all nodes of the cluster:

CentOS and RHEL:

```
yum install smartsense-hst -y
```

Ubuntu and Debian:

```
apt-get update  
apt-get install smartsense-hst -y
```

SUSE:

```
zypper install smartsense-hst -y
```



**Note:**

You must perform this step on all nodes of the cluster.

3. Log into the node which runs ambari server.
4. Run the following command to add Smartsense to the HDF stack:

```
hst add-to-ambari
```

5. Restart the ambari server:

```
ambari-server restart
```

After performing these steps, log in to the Ambari web UI and add the SmartSense by using the "Add Service" wizard.

**Related Information**

[Add the SmartSense service](#)

## HST server placement

You should designate one node in the HDP cluster as the HST server, so that this component can efficiently consolidate the data collected by all HST agents into a single downloadable file (referred to as a *bundle*). Any of the management nodes, such as Ambari Server, Metrics Server, and so on, are good choices for the HST server placement.

Administrators and each HST agent in the cluster must have network access to the HST server. This connectivity is required for agents to consolidate their data and for Hadoop administrators to download completed bundles. For a full list of ports and a data flow diagram, refer to SmartSense ports and traffic flow documentation.

**Related Information**

[SmartSense ports and traffic flow](#)

## Add the SmartSense service

Use these steps to add SmartSense service in Ambari.

**Before you begin**

- You should know your SmartSense ID and account name (both are available in the Hortonworks support portal in the **Tools** tab).
- You must also ensure that an Ambari agent is running on the same host as the Ambari server.

**Procedure**

1. From the Ambari web UI, select **Add Service** from the ... drop-down menu.
2. From the list of installable services, select **SmartSense**, and then click **Next**.

3. On the **Assign Masters** page, select cluster node for the HST server and then click **Next**.
  - For a list of criteria to determine the best node to select for HST server, see the HST server placement documentation.
4. On the **Customize Services** page, validate the values in the following fields, as appropriate to your environment:

**Table 1: Ambari web UI options for customizing SmartSense service**

UI Option	Description
Configuration Tab: Basic Property: Customer account name	Your account name, available from the <b>Tools</b> tab in Hortonworks support portal
Configuration Tab: Basic Property: SmartSense ID	Your SmartSense ID, available from the <b>Tools</b> tab in Hortonworks support portal
Configuration Tab: Basic Property: Notification Email	The email address notified when SmartSense bundles have been received and recommendations are ready for your review
Configuration Tab: Basic Property: Enable Flex Subscription	Use this option only if you have an existing Hortonworks flex support subscription. You must enter your Flex Subscription ID.
Configuration Tab: Basic Property: Bundle Storage Directory	The directory on the HST server that will be used to store completed bundles  Because bundles can be large, this directory should have at least 1GB of free space.
Configuration Tab: Basic Property: Server Temporary Data Directory	The directory on the HST server that is used to assemble results from HST Agents into completed bundles  This directory must be large enough to handle the intermediate results of HST agent collection data: at least 5 GB of free space.
Configuration Tab: Basic Property: Agent Temporary Data Directory	The directory on the HST server that is used to assemble results from HST Agents into completed bundles  This directory must be large enough to handle the intermediate results of HST agent collection data: at least 5 GB of free space.

Click **Next**.

The Ambari Stack Advisor assesses your cluster configuration and might alert you to configuration issues. Note that this is not related to SmartSense, and is simply what Ambari does upon adding any service. SmartSense never makes configuration changes to your cluster. No cluster services need to be restarted after installing SmartSense, and any configuration changes that are noticed should be reverted.

If you have a kerberized cluster, you will be prompted for the KDC admin credentials during this step.

5. On the **Review** page, click **Deploy** to complete your SmartSense service installation.
6. Exit the wizard.

### Related Information

[HST server placement](#)

## Installing SmartSense gateway

If your HST server host does not have outbound internet access, you can connect it to a single SmartSense gateway that does. A single gateway can support multiple internal HST server deployments, uploading all of their bundles to Hortonworks for support as well as proactive analysis.

Using the SmartSense gateway involves knowing where to place it, installing it, and integrating it with Ambari.

## SmartSense gateway placement

You must deploy the SmartSense gateway on a host in a network zone that has both outbound internet access and inbound access from all HST server instances. The connectivity between the HST server and the gateway is secured using mutually authenticated SSL.

By default, TCP port 9450 is used to register HST server instances with the gateway. After initial registration, TCP port 9451 is used for the authenticated API communication between the HST server and the gateway. Both the registration and API communication ports are configurable in the gateway `hst-gateway.ini` file.

Outbound connections from the SmartSense gateway to the external Hortonworks SmartSense analysis environment use HTTPS to transmit bundles to Hortonworks. Specific connectivity details for the SmartSense environment are outlined in [SmartSense ports and traffic flow](#).

### Related Information

[SmartSense ports and traffic flow](#)

## Install and start SmartSense gateway

The SmartSense gateway is not managed by Ambari and must be installed manually. It is included in the same `smartsense-hst-$HST_VERSION` package used for the HST server and HST agent. You can access this package in your Ambari repository; Additionally, it is available on the Tools tab of the Hortonworks support portal. Once you have the package, perform these steps:

### Procedure

1. Install the SmartSense package on the chosen gateway host:

- RHEL, CentOS, or SLES:

```
# rpm -ivh smartsense-hst-$HST_VERSION.rpm
```

- Ubuntu or Debian:

```
# dpkg -i smartsense-hst_$HST_VERSION.deb
```

2. Configure the gateway by editing the `/etc/hst/conf/hst-gateway.ini` file:

- Specify the path to the JDK to be used by the gateway in the `[java]` section:

```
[ java ]
home={ $path_to_your_JAVA_HOME }
```

- To configure HTTPS connectivity, refer to [SmartSense gateway setup](#) for the HTTPS user name, password, host, and port details. To view this article, you need a valid Hortonworks support account.

3. Start the gateway:

```
# hst gateway start
```

On startup, the gateway attempts to connect to the configured HTTPS host, and if the connection fails, the gateway does not start and logs the reason for connectivity failure to the `/var/log/hst/hst-gateway.log` file.

4. Integrate the gateway with the HST server by following the instructions for integrating the gateway with Ambari-managed SmartSense .

### What to do next

If you deploy the gateway on a server that is also hosting an HST agent, and that HST agent has been deployed through Ambari, and Ambari is configured for nonroot operation, you should run the SmartSense gateway as the same user that the Ambari agent is configured to run as.

Additionally in this specific scenario, you must modify the following permissions using the commands below. In this example "ambari" is the user the Ambari agent has been configured to run as:

```
chown -R ambari:hadoop /var/lib/smartsense/hst-gateway
chown -R ambari:hadoop /var/log/hst
chown -R ambari:hadoop /var/run/hst
```

### Related Information

[Integrating gateway with Ambari-managed SmartSense](#)

[SmartSense Gateway Setup \(Salesforce Article\)](#)

## Integrating gateway with Ambari-managed SmartSense

Since SmartSense gateway is embedded, there is no need to preform these additional configuration steps to integrate the gateway.

You only need to perform them if your gateway server is not embedded (i.e. explicitly installed on some host other than HST server).

In such a case, you need to specify the gateway host in the Gateway configuration tab in Ambari 2.4.x or later, by providing the fully qualified domain name of the host running the gateway.

**Table 2: Ambari web UI options for integrating HST gateway**

Ambari 2.4+	Note
Configuration Tab: Gateway Property: Gateway host	The fully qualified domain name of the host running the gateway

When enabled, the gateway automatically uploads completed bundles to Hortonworks when a capture is completed.

## SmartSense uninstallation

Use these instructions to uninstall SmartSense gateway.

### Uninstall SmartSense Gateway

To remove the SmartSense gateway, follow these steps:

#### Procedure

1. Ensure that the SmartSense gateway is stopped:

```
# hst gateway stop
```

2. Remove the smartsense-hst package:

- RHEL, CentOS, r SLES:

```
# rpm -e smartsense-hst
```

- Ubuntu or Debian:

```
# dpkg -r smartsense-hst
```

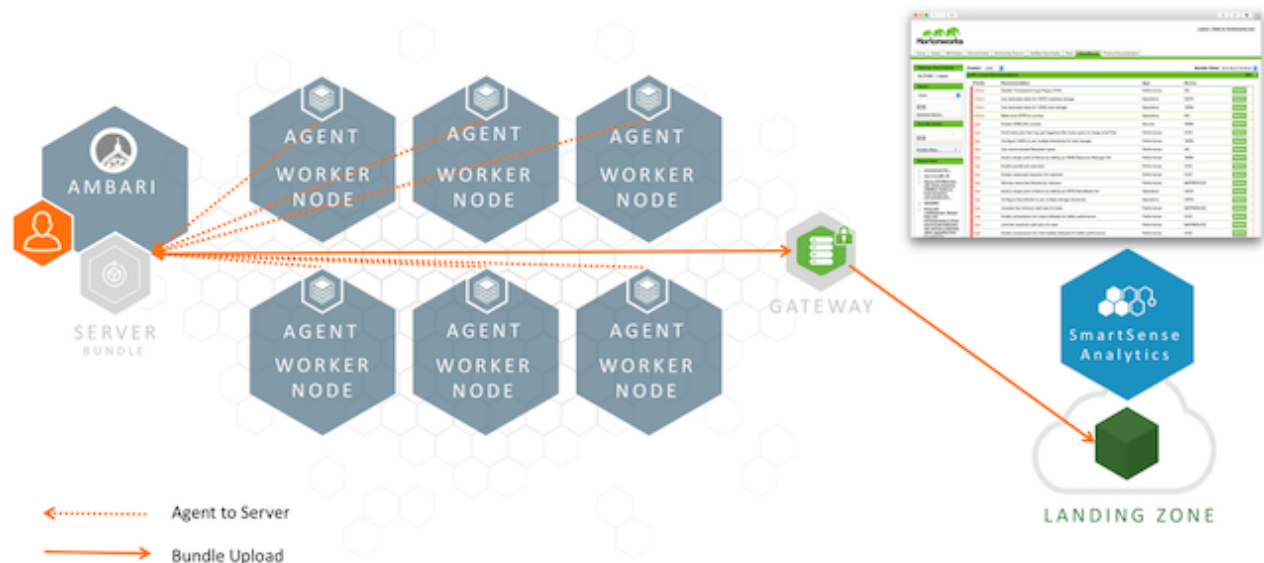
### 3. Remove logs produced by the gateway:

```
# rm /var/log/hst/hst-gateway.*
```

## SmartSense ports and traffic flow

When deploying SmartSense in an enterprise environment, you must carefully plan your network architecture. SmartSense functionality relies on how multiple internal clusters create bundles and then send them through a central SmartSense gateway to the hosted Hortonworks environment for analysis or to Hortonworks support to troubleshoot support cases.

The following graphic illustrates how data traffic flows along various communication channels between cluster nodes and SmartSense ports:



The communication channels illustrated here are described in the following sections:

### User interface or Ambari View to HST server

When using SmartSense without Ambari, users access the web UI directly, whereas when using Ambari, they use Ambari View to communicate with the server.

**Table 3: Ports: User interface or Ambari View to HST server**

Source Component	Destination Component	Destination Port	Purpose
User interface or Ambari View	HST server	tcp/9000	Web UI communication
Zeppelin activity explorer UI	Zeppelin server	tcp/9060	Web UI communication

### HST agents to HST server

All communication between the HST server and HST agents is initiated by the agent, using the following ports:



**Table 4: Ports: HST agents to HST server**

Source Component	Destination Component	Destination Port	Transport Security	Purpose
HST agent	HST server	tcp/9440	One-way SSL	Agent registration
HST agent	HST server	tcp/9441	Two-way SSL	Anonymized bundle transfer

HST agents register themselves with the HST server, and when invoked to capture data, use the same port to securely transmit captured data back to the HST server.

## HST server to SmartSense gateway

All communication between the HST server and the SmartSense gateway is initiated by the server, using the following ports:

**Table 5: Ports: HST server to SmartSense gateway**

Source Component	Destination Component	Destination Port	Transport Security	Purpose
HST server	SmartSense gateway	tcp/9450	One-way SSL	HST server registration
HST server	SmartSense gateway	tcp/9451	Two-way SSL	Encrypted bundle transfer

HST servers register themselves with the SmartSense gateway using the two-way SSL registration port (tcp/9451), and when bundle capture is complete, this port is used to securely stream the bundle file to the SmartSense gateway.

## SmartSense gateway to Hortonworks

Hortonworks does not initiate communications to the SmartSense gateway, all communication is initiated by the SmartSense gateway to Hortonworks. For this interaction, the following ports are used:

**Table 6: SmartSense Gateway to Hortonworks**

Source Component	Destination Component	Destination Port	Purpose
Gateway	Hortonworks	tcp/443	HTTPS bundle upload

As bundles are captured, the HST server uses the two-way SSL communication channel to securely stream the bundle file to the SmartSense gateway. After this process starts, the SmartSense gateway opens up a secure connection to Hortonworks, using the HTTPS port to upload the bundle.

You have two options when configuring the communication between the SmartSense Gateway and Hortonworks:

- Allow firewall access from the gateway to a CNAME using port 443.

The Hortonworks HTTPS servers utilize Elastic Load Balancing from Amazon Web Services. The CNAME is recommended as the number of instances, and IPs of instances used by the load balancer are fluid. Using the CNAME provides the greatest availability.

- Allow firewall access from the gateway to a pair of static IPs using port 443.

These IPs do not change, and they use "round-robin" DNS for load balancing. This is the least preferred option, because instance availability is not quickly updated in DNS.

For Details about setting up the SmartSense gateway see our [SmartSense gateway setup](#) Knowledge Base article. To get the specific details on the static IP addresses used with SmartSense, see our [SmartSense Static IPs](#) Knowledge Base article. To view both of these articles, you need a valid Hortonworks support account.

### Related Information

[SmartSense Gateway Setup \(Salesforce Article\)](#)

## SmartSense installation troubleshooting

This section includes steps for troubleshooting issues that you might encounter during SmartSense installation.

### SmartSense SSL troubleshooting

SmartSense components use SSL for protecting communications between the HST server and agents, and between the HST server and SmartSense gateway. If installation issues arise, you can reset these SSL certificates.

#### HST Server

1. To reset the HST server SSL certificate database, which forces all HST agents to regenerate their certificates, use the `hst reset` command:

```
# hst reset
Resetting SmartSense Server will remove server and all registered agent
certificates and reset the certificate database. Do you want to continue?
[y/n] (default: n): y
SmartSense Server is currently running and needs to be stopped in order to
reset. Do you want to stop the SmartSense Server? [y/n] (default: n): y
SmartSense Server stopped
SmartSense Server reset completed.
Do you want to restart SmartSense Server? [y/n] (default: y): y
Server PID at: /var/run/hst/hst-server.pid
Server out at: /var/log/hst/hst-server.out
Server log at: /var/log/hst/hst-serer.log
Waiting for server start . . . . .
```

2. Next, you must manually reset each individual HST agent after running this command. For instructions on how to reset the agents, see the following HST Agent section.

#### HST agent

Perform these steps in the following cases:

- An individual agent is having issues related to SSL when communicating with the HST server.
- You have just reset the HST server SSL certificate database (see the HST Server section above). In this case, you must perform these steps on each individual HST agent.

1. Use the `hst reset-agent` command to remove all certificates registered with the HST server for the specific agent.
2. Next, run `hst setup-agent -q` to register the agent with the server and download new certificates.

```
# hst reset-agent
Resetting SmartSense Agent will remove all certificates registered with
SmartSense server. Do you want to continue? [y/n] (default: n): y
SmartSense Agent reset completed.
# hst setup-agent -q
```

#### SmartSense gateway

If HST server is having issues related to SSL when communicating with the SmartSense gateway, you can use the `hst gateway reset` to remove all HST server certificates registered with the specific gateway.

From the SmartSense gateway, you can execute the following process:

```
# hst gateway reset
Resetting SmartSense Gateway will remove all certificates and reset the
certificate database. Do you want to continue? [y/n] (default: n): y
SmartSense Gateway stopped
```

```
SmartSense Gateway reset completed.
Gateway has to be started to create new certificates. Do you want to start
the Gateway? [y/n] (default: y): y
SmartSense Gateway PID at: /var/run/hst/hst-gateway.pid
SmartSense Gateway out at: /var/log/hst/hst-gateway.out
SmartSense Gateway log at: /var/log/hst/hst-gateway.log
Waiting for Gateway start . . . . .
SmartSense Gateway started.
```

## Reporting issues

If you have encountered a functional issue or observed a security issue, you can raise a support ticket in the [Hortonworks support portal](#). To open a new support case, navigate to the **Cases** tab and click **Create New Case**. During case creation choose **Product Component: SmartSense**.

### Related Information

[Hortonworks Support Portal](#)

## Enabling flex support subscription

If you would like to use an existing Hortonworks flex support subscription for your cluster, you can enable it during SmartSense service installation or later by using the following steps.

### Procedure

1. From the Ambari web UI, select the SmartSense service.
2. Select Configs>Basic.
3. Click the toggle button next to Enable Flex Subscription to enable flex subscription.
4. Enter your Flex Subscription ID, for example, "FLEX-01234567889".
5. Click the Save button to save the configuration changes.
6. Enter a description for the configuration change and click Save.
7. Click OK to confirm.
8. Click Restart>Restart All Affected to restart SmartSense service.

### Results

After performing these steps, your flex subscription is enabled for the cluster.

### Related Information

[Hortonworks Flex Support Subscription](#)